



Data-centric Networking for a Data-centric IoT: A User's Perspective

Eve M. Schooler

Technology, Strategy, Pathfinding
Internet of Things Group (IoTG)

May 31, 2016



Outline

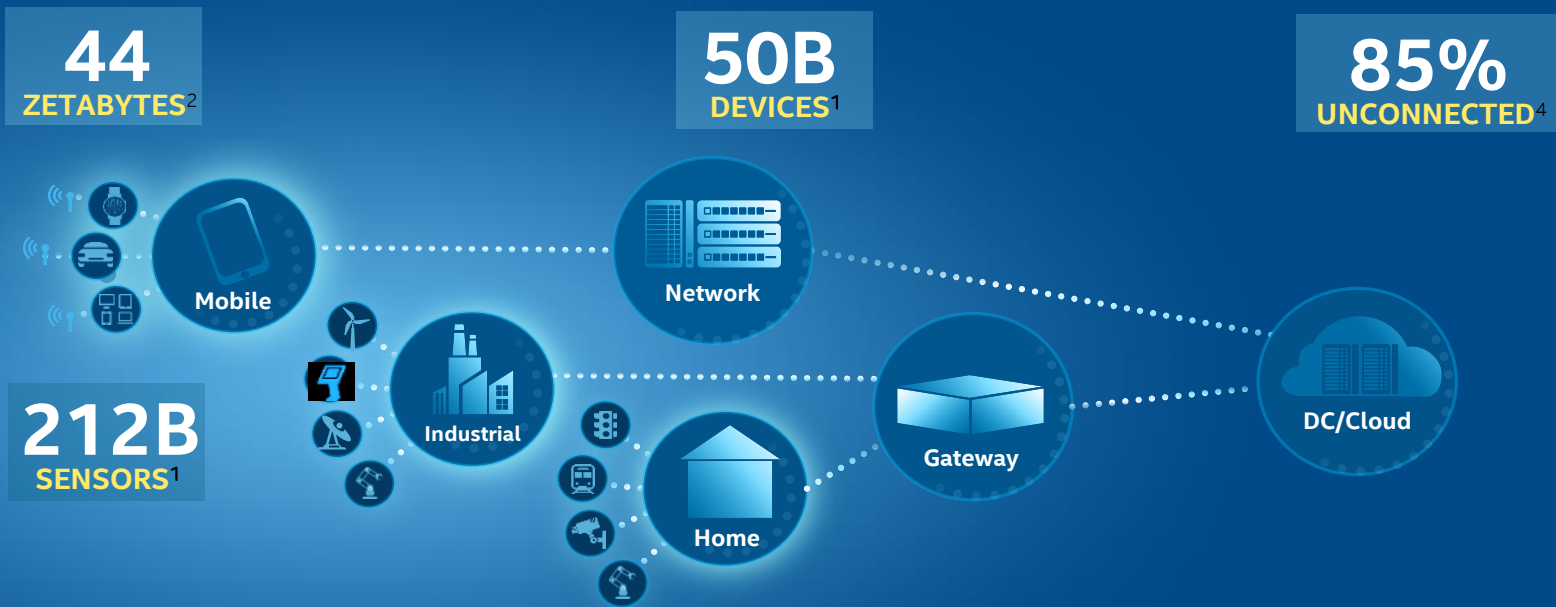
- Context
- What makes IoT interesting...and disruptive?
- Why is ICN well-suited for IoT?
- IoT Use Cases – A Sampling of ICN benefits
 - The Smart Grid – Smart Home, Smart Buildings, Smart Neighborhoods
 - Trusted Analytics at the Network Edge - Critical Infrastructure, Remote Monitoring and Interactive Control (Manufacturing, Transportation, etc)
 - Interoperability - Smart Objects and Trust
- Lessons learned, Gaps, Call to action!

Acknowledgment

This work is the result of many collaborators (where they were when the collaboration began/where they are currently and if they are **currently at Intel**):

Moreno Ambrosin (U. Padua/Intel IoTG), **Andrew Brown** (Intel IoTG), **David Cohen** (Intel DCG), **Mihaela Ion** (U. Trento/Google), **Sanjana Kamath** (Intel IoTG), **Sung Lee** (Intel IoTG), **Qinghua Li** (Penn State/U. Arkansas), **Hassnaa Moustafa** (Intel IoTG), **Adedamola Omotosho** (Intel QSD), **Sebastian Schoenberg** (Intel IoTG), **Matthias Schunter** (Intel Labs), **Jeff Sedayao** (Intel SW & Services Group), **Karen Sollins** (MIT), **Xinlei Wang** (UC Davis/Facebook), **Dave Zage** (Purdue/Intel IoTG), **Jianqing Zhang** (Intel Labs/Vmware).

The IoT ...



COST OF SENSORS PAST 10 YEARS ¹	2X ↓	COST OF BANDWIDTH PAST 10 YEARS ²	40X ↓	COST OF PROCESSING PAST 10 YEARS ³	60X ↓	COST OF STORAGE PAST 10 YEARS ^{5,6}	25X ↓
--	----------------	--	-----------------	---	-----------------	--	-----------------

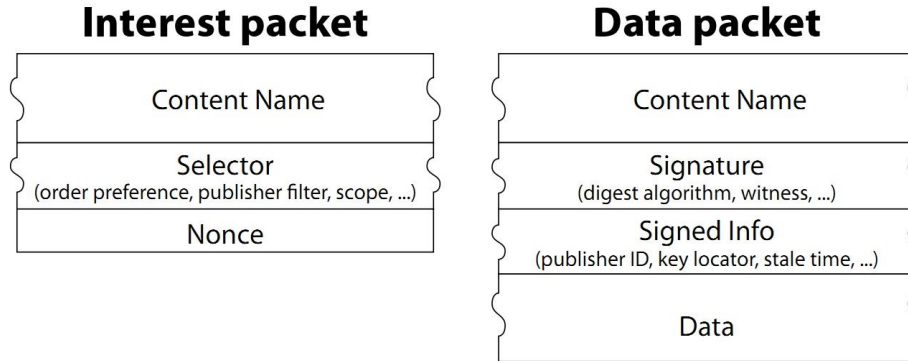
...changes our architectural thinking



1. IDC
2. MC/EDC: The Digital Universe of Opportunities
3. Goldman Sachs
4. IMS Research
5. Cost per Gigabyte Update
6. Gartner

Attractive ICN Properties

- Name-based data routing
- Distributed data caching
- Self-contained data security



IoT Use Cases – A Sampling of ICN Benefits

Smart Grid

- **iHEMS:** ICN-based Home Energy Management System (HEMS)
- **iCity:** ICN-based Neighborhood-Coordinated Electric Vehicle (EV) Charging
 - Constrained (Mobile) Device Optimizations

Trusted Analytics at the Network Edge

- **Updicator:** ICN-based SW updates to O(Billions) of devices
- **Remote monitoring and interactive control**
 - Video use cases for IoT (Manufacturing, Transportation, etc)
 - Fog Computing and Smart Data “Pipes”

Interoperability

- **Smart Objects**

Perspectives from the Smart Grid

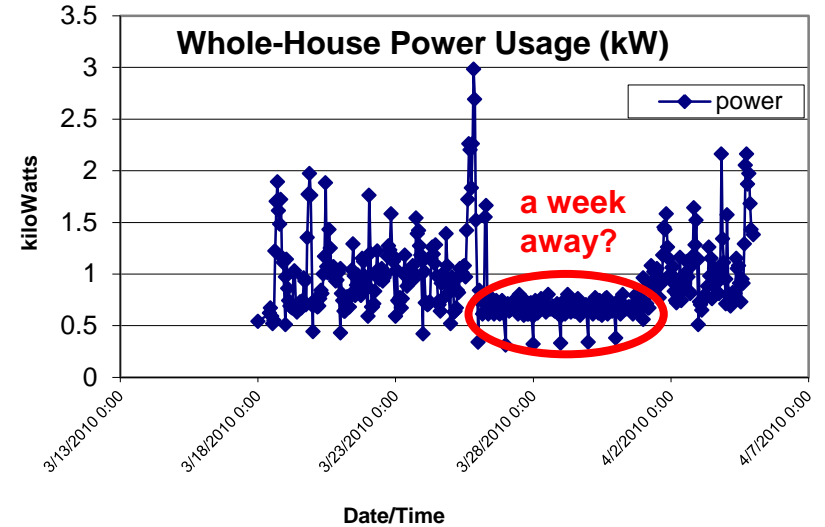
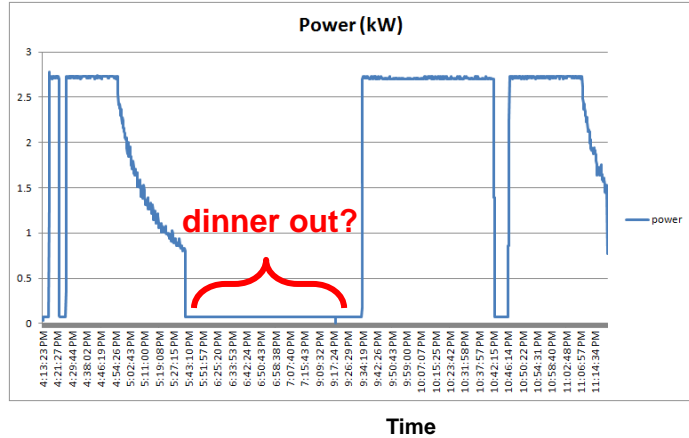
iHEMS: ICN-based Home Energy Management System Challenges

- Proliferation of smart “devices”, many mobile
 - Many in unmanaged or self-managed nets
 - Average users (vs. IT experts) “managing” them
- Tidal wave of data generated
 - O(Petabytes) in Smart Grid data alone
- **Security & Privacy increasingly critical**
 - IoT data can reveal personal identities, behaviors, location, health, etc.
 - Scalable & flexible data encryption required




Energy Data & Privacy

Electric Vehicle Charging (kW)




Nocturnal Device Power Usage (kW)

CPAP machine
for Sleep Apnea





Entertainment Device Usage (kW)



iHEMS: ICN-based Home Energy Management System

Opportunities

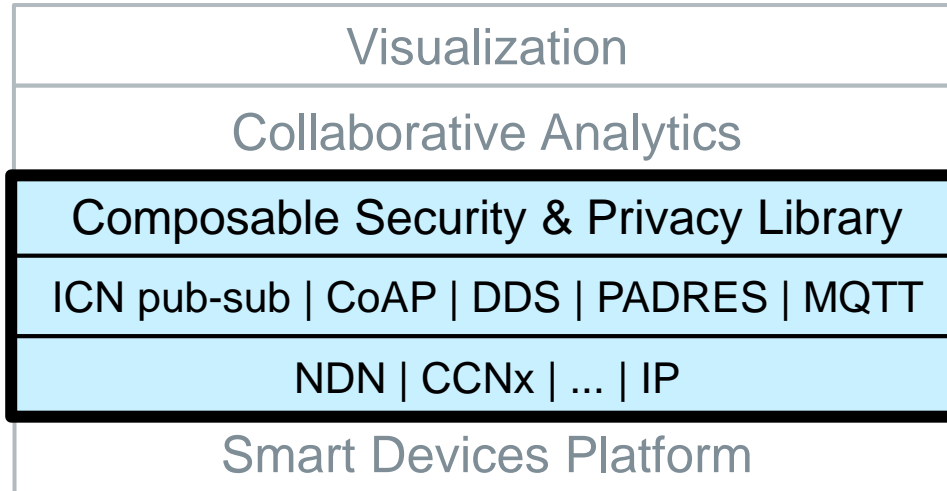
- Avoid replacing “thin waist” of the entire Internet, instead...
- Enable ICN-based Trusted Local Clouds for IoT at the Network Edge – *akin to uGrids in the Smart Grid*



Define the Anatomy of the Trusted Data Bus

Opportunities

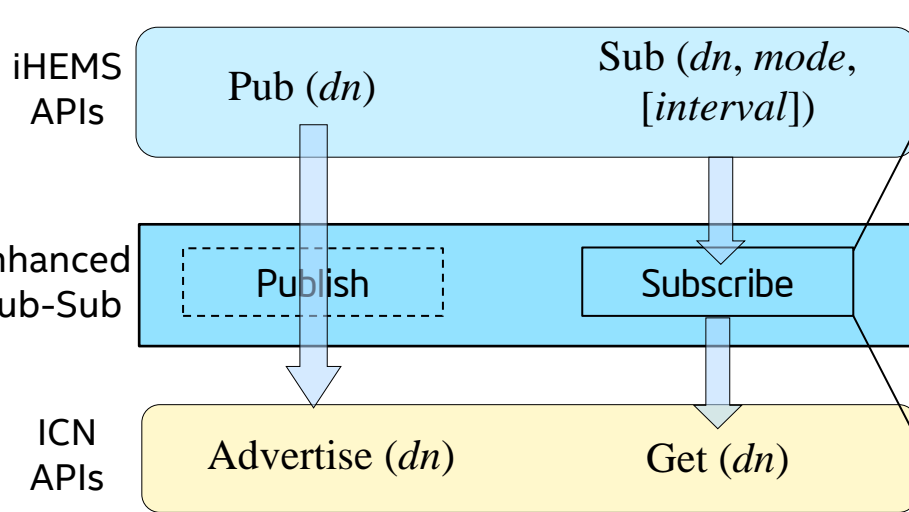
- Re-usable, interchangeable ICNs and middleware
 - Plug-in-play SW modules



- Composable uServices

Data integrity
Data authenticity
Data confidentiality
Subscription confidentiality
Scalable key management
Encrypted filtering
Publisher/Subscriber anonymity
...


Use-case Driven Requirements and APIs



Type	Feature	Example
Instant	Subscribe <i>once</i> Deliver <i>once</i> immediately	Quick check of power usage
Persistent	Subscribe <i>once</i> Deliver <i>multiple times</i> , upon event	Power events notification
Periodic	Subscribe <i>once</i> Deliver <i>periodically</i> at specified interval	Sampled data like temperature, voltage, etc.
Constrained	Subscribe <i>once</i> Deliver <i>multiple times</i> , if condition met	Alarms on temperature, power, etc.

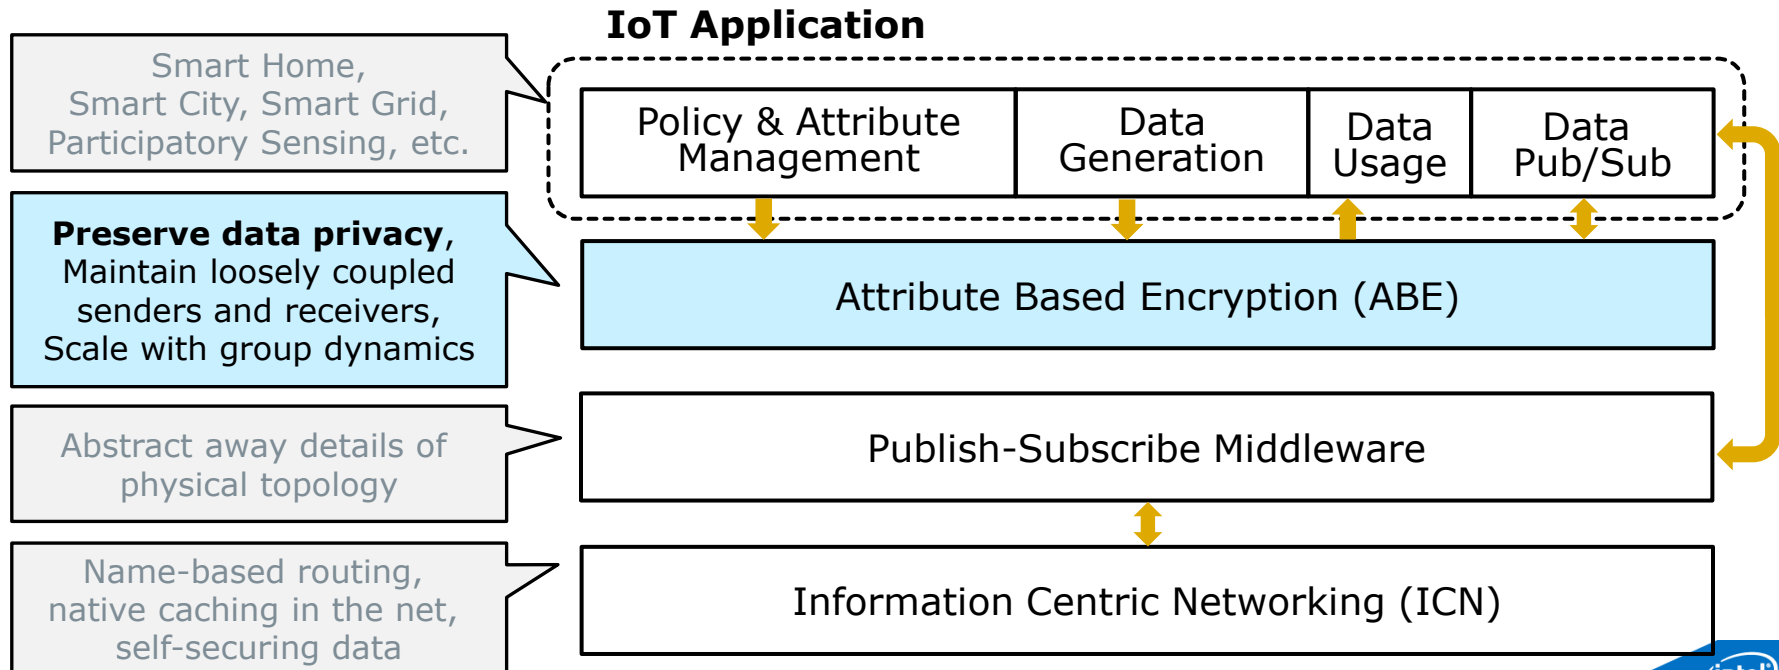
iCity: Neighborhood-Coordinated EV Charging Challenges

- Not all EVs can charge at once
 - Risk damage to transformers
- 100x to 1000x more devices and sensors
- Dynamics of mobility
- Wider geographic distributions
 - Data flows across private-public boundaries
- **Scalability of Data Privacy solution**



iCity: Neighborhood-coordinated EV Charging Opportunity

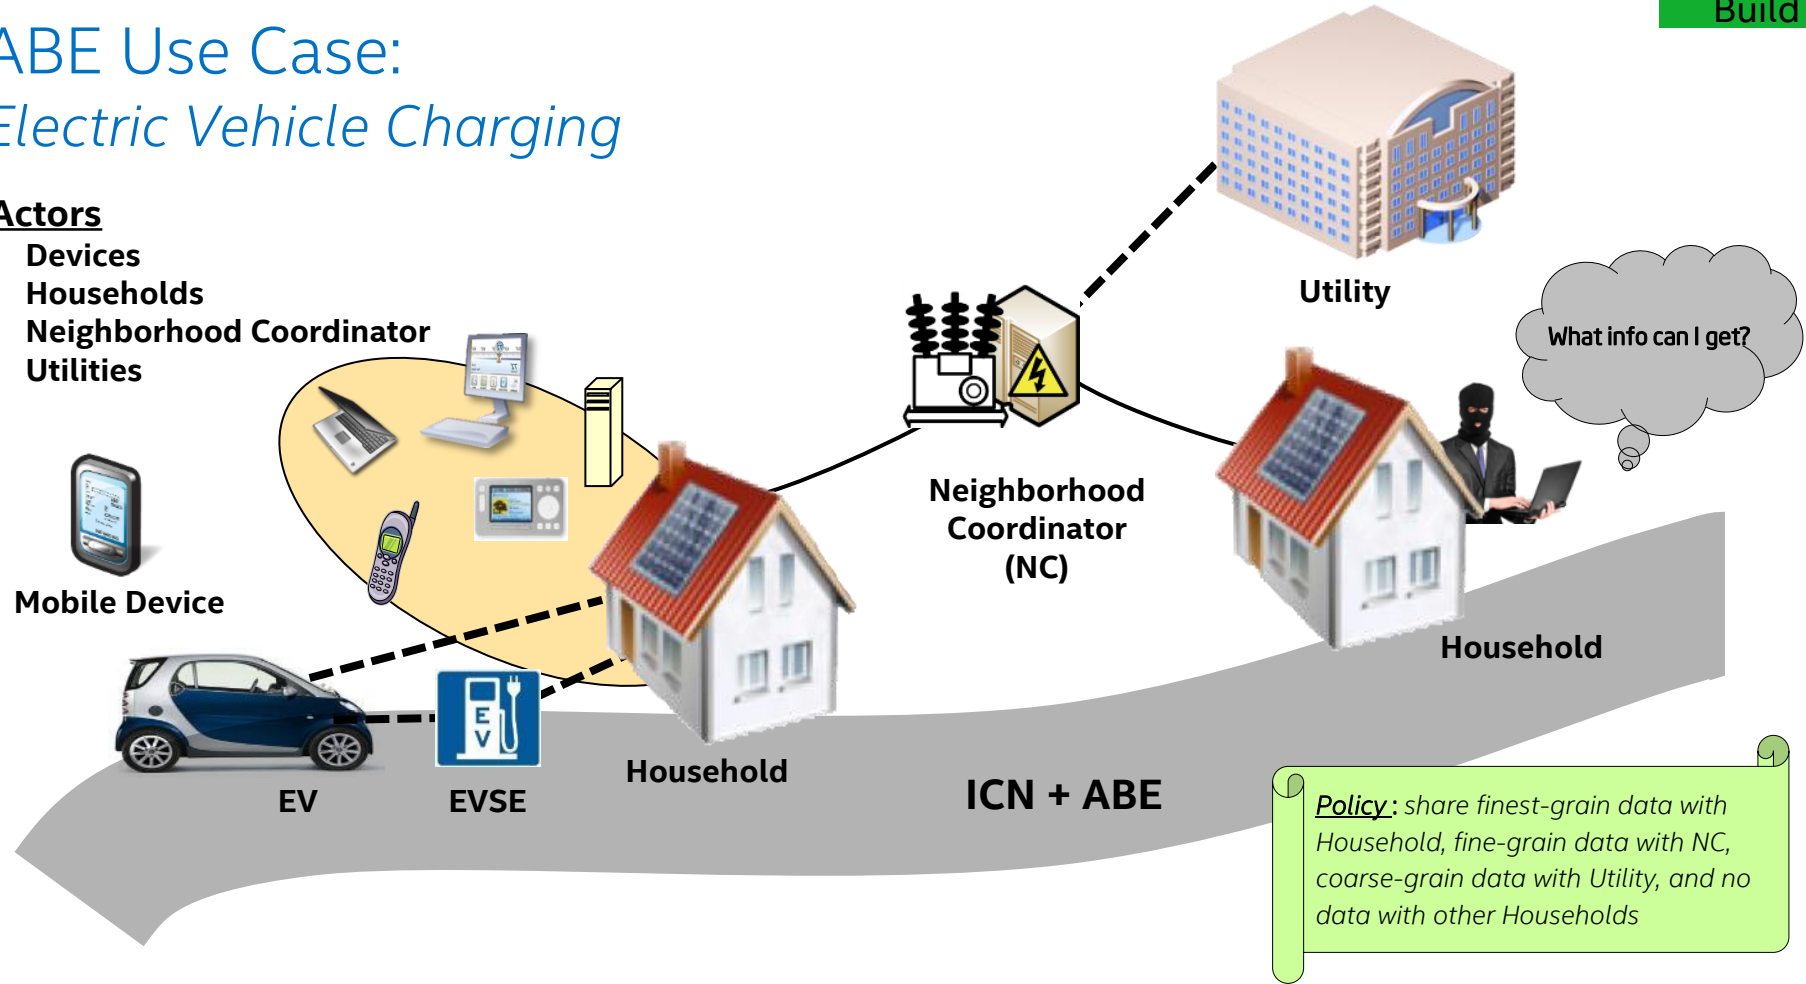
- Revisit Trusted Data Bus - and extend with Data-centric Privacy



ABE Use Case: Electric Vehicle Charging

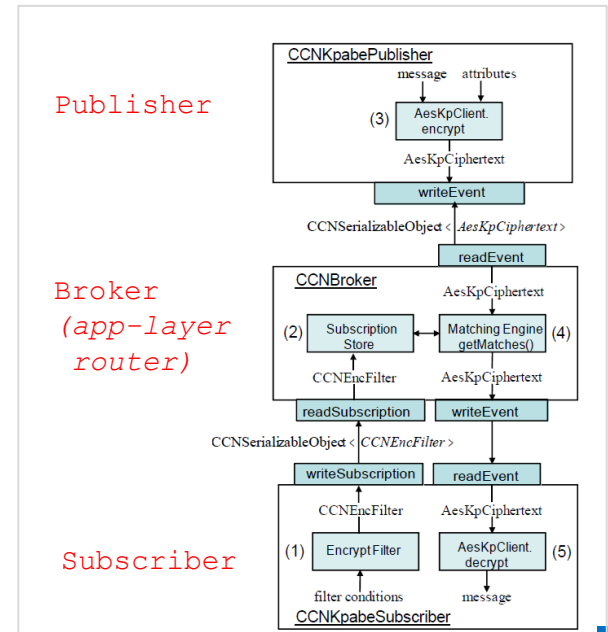
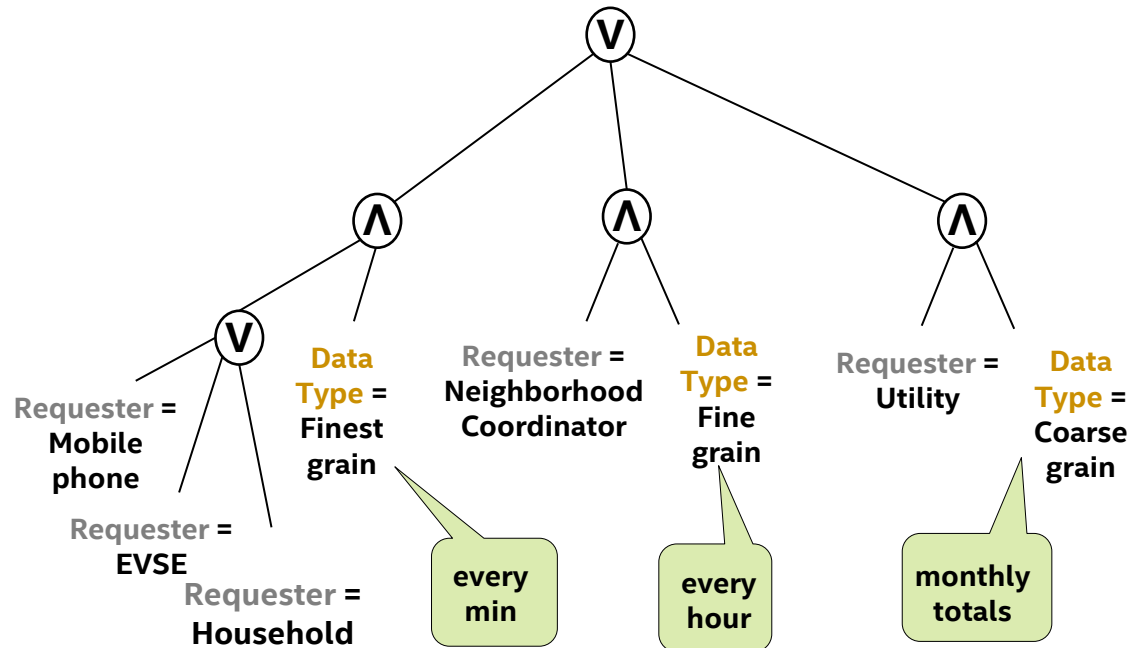
Actors

- Devices
- Households
- Neighborhood Coordinator
- Utilities



iCity: Neighborhood-Coordinated EV Charging Opportunity

- Attribute-based Policies are embedded ...and stay with the Data
- Richer attribute-based pub-sub interface at the application layer

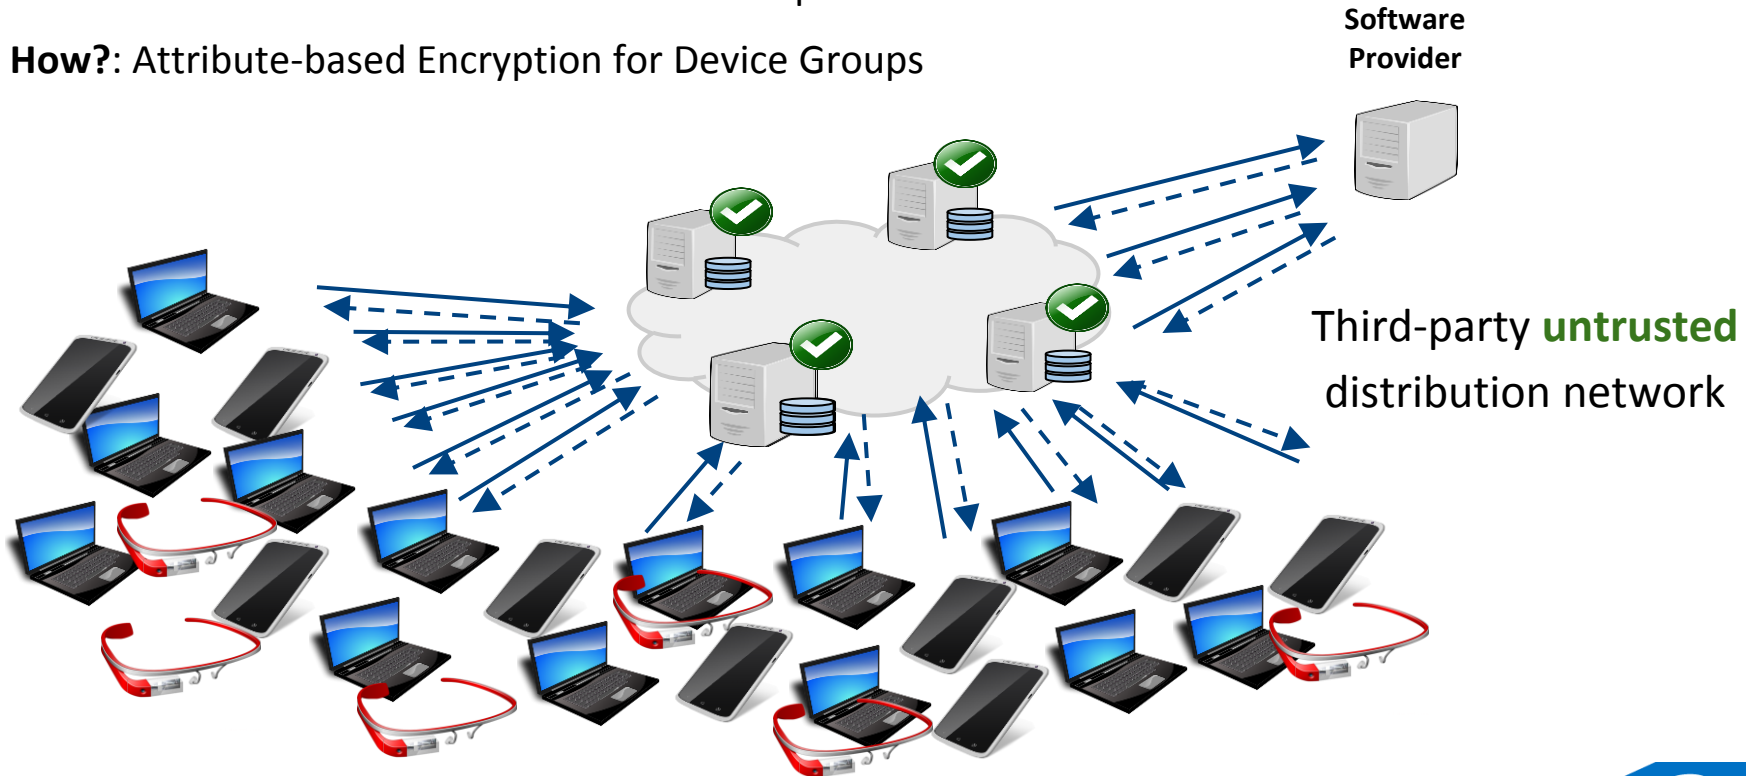


Toward Trusted Analytics at the Network Edge

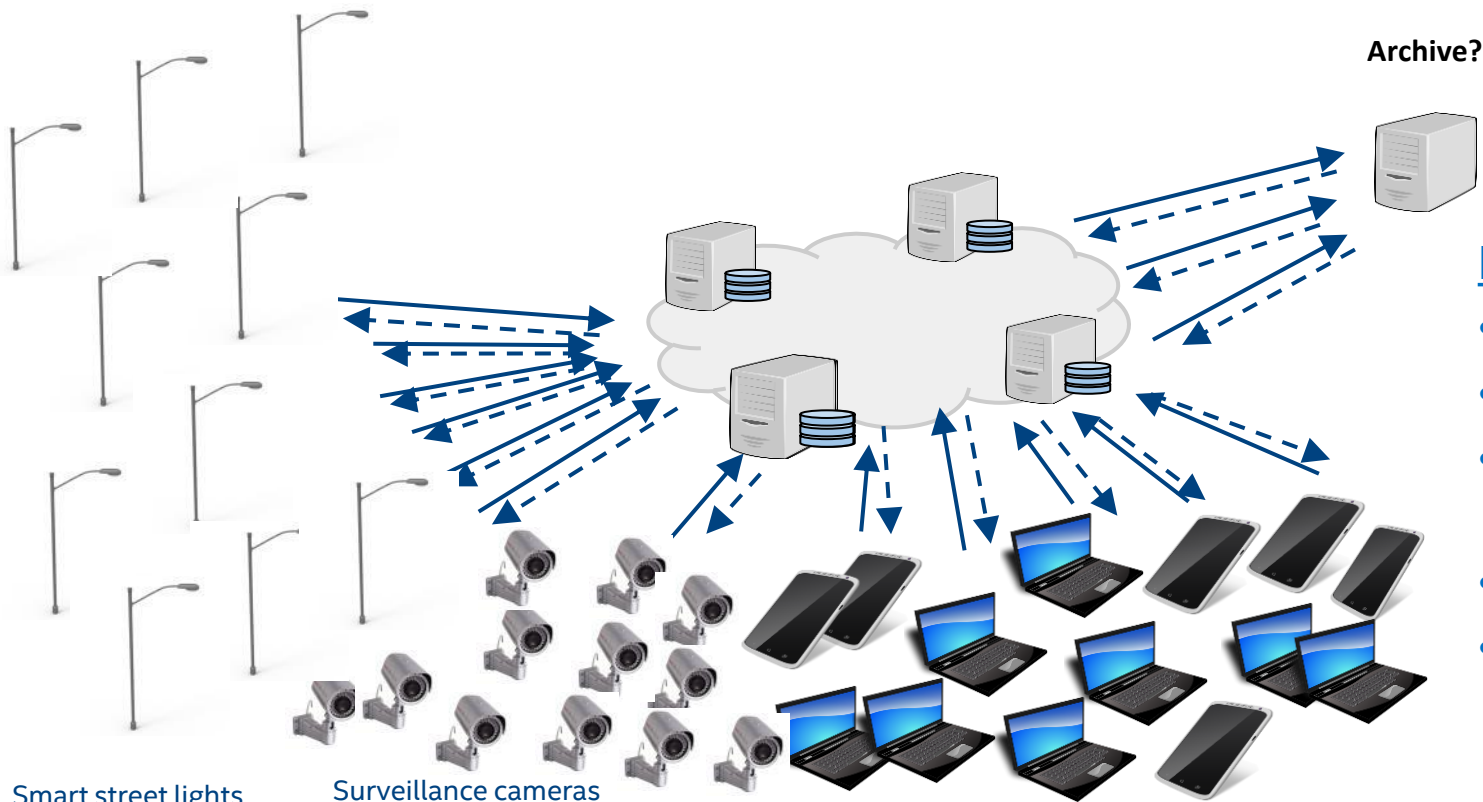
Updicator: Updating Billions of IoT Devices

What?: Use untrusted NDN to Scale Secure Updates

How?: Attribute-based Encryption for Device Groups



Remote Monitoring and Interactive Control: *What if all Things always streamed real-time (video) data?*



Smart street lights

Surveillance cameras

Many Use Cases


- Manufacturing
- Smart Grid
- Building Surveillance
- Transportation
- Healthcare & Eldercare

Data Inversion Problem: IoT data originates at the “Edge”

Result: Cloud functionality migrating to be more proximate to the data

“By 2018, 40% of IoT-created data will be stored, processed, analyzed, and acted upon closest to, or at the edge of the network.”

12/2014



Problem: Legacy clouds fall short ...or are unusable

When the IoT data generated is

- Delay-sensitive
- High-volume
- Trust-sensitive
- (Intermittently) Disconnected


Countless examples

- Both near term & further out

Video Analytics



Augmented Reality



- Data heavy
- Compute intensive
- **Response times <30ms**
- Small form factor
- Low power

Use ICN to move the “compute” (executables) to the “data” (observations) at or nearer the network edge

Remote Monitoring & Interactive Control

Challenges

- Where to perform analytics?
 - Tandem need for analytics on the analytics?
- Timeliness...
 - ...of video data delivery?
 - ...of response and/or action?
 - **How timely is timely enough?**
- Interactivity & control loops

Opportunities

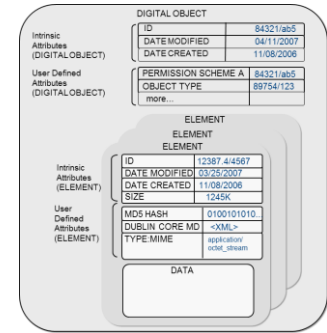
- ICN for (near) real-time Fog
 - Re-usable comms/storage
- ICN for generalized anomaly detection/normalcy baselining
- Trusted data bus evolution
 - ICN for Real-time (WebRTC, TSN-aware)
- Smart Data “Pipe”?
 - Reverse CDN and ICN
 - Named data, tagged data (interesting events, features)

Musings on Interoperability

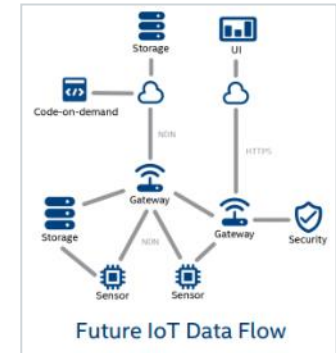
Smart Objects

Bridging the Cyber-Physical Divide

- Every object has a unique DO identifier (DOI)
- Objects can be: devices, users, comms channels, data, meta-data, clouds, services, algorithms, etc.
- Smart Objects (SOs) = Self-describing Things
 - Enable data interoperability, microservice composition, policy management, smart data pipes, etc.
- Widespread discussion
 - OpenFog, OCF, IPSO, IETF, NIST, ITU, OPC, etc.
 - Taxonomies, Registries, Bridges
- Action: Expose Data Plane **meta-data** to seed SOs



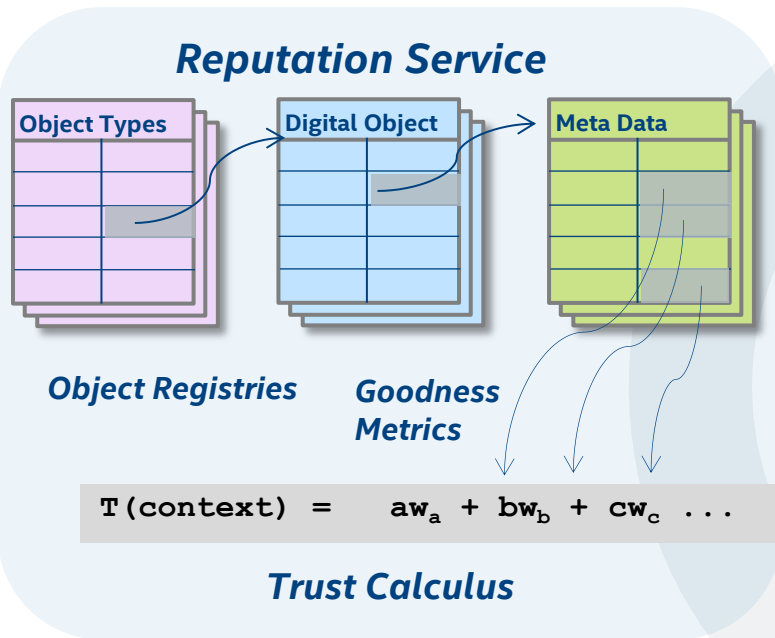
E.g., ITU/IETF Digital Object (DO) Architecture



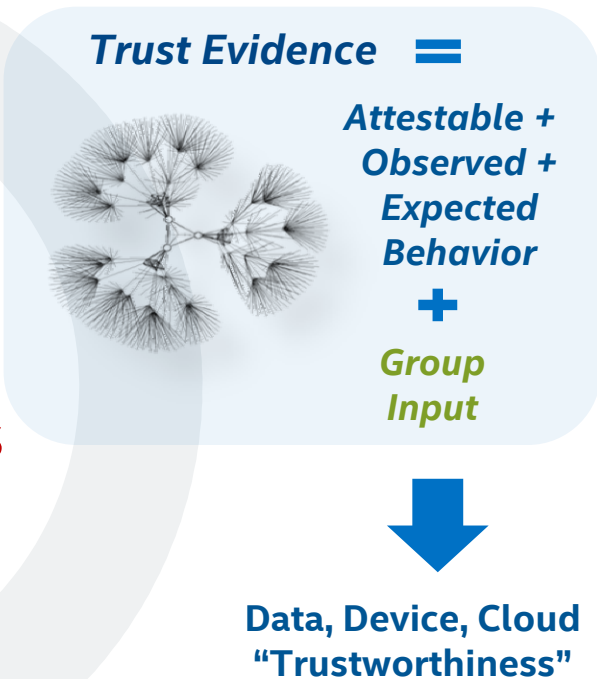

[1] NDNComm'15

How to relate upper-layer Objects with lower-layer ICN names?

Toward an Attribute-based Trust Framework: *Building Trust and Trust Anchors*



**Ubiquitous
Monitoring,
Measurement,
Tiered Analytics**



Visualize Reputation

A Personal (Marauder's) Map to:

- Organize the sea of data
- Map cyber data to physical spaces
- Make the invisible visible
- Visualize device relationships/reputations
- Establish easy-to-setup privacy policies
- Disallow “spoofing”

Decide which of the 212B sensors, 50B devices, and 100K clouds...

- to connect to
- to allow connections from



A Call to Action:

How to make ICN even more successful?!

- One voice in the standards community
- Privacy and Caching – friends or foe?
 - Encryption everywhere, all the time
- Quantify/Qualify ICN benefits
 - Up-the-stack and down-the-stack

Questions?


Back-up

References

1. Andrew Brown, Sebastian Schoenberg, Eve Schooler, “NDN and the Internet of Things: Analytics Everywhere”, poster, *2nd Annual Named-Data Networking Community Meeting*, NDNComm’15, LA, CA (Sept 2015).
2. Andrew Brown et al, “Information Centric Networking for IoT Devices”, *Intel Software Professionals Conference*, demo & presentation (Aug 2015).
3. David E. Cohen and Eve M. Schooler, “Data Inversion and SDN Peering: Harbingers of Edge Cloud Migration”, *IEEE ComSoc MMTC E-letter*, Special issue on Big Data in 5G Networks, Vol.9, No.6 (Nov 2014).
4. Moreno Ambrosin, Christoph Busold, Mauro Conti, Ahmad-Reza Sadeghi, Matthias Schunter, “Updicator: Updating Billions of Devices by an Efficient, Scalable and Secure Software Update Distribution Over Untrusted Cache-enabled Networks”, *ESORICS’14* (Sept 2014).
5. Xinlei Wang, Jianqing Zhang, Eve M. Schooler, “Performance Evaluation of Attribute-based Encryption: Toward Privacy in the IoT”, *IEEE ICC’14*, Sydney, Australia (Jun 2014).
6. Mihaela Ion, Jianqing Zhang, Eve M. Schooler, “Toward Content-Centric Privacy in ICN: Attribute-based Encryption and Routing”, *ACM SIGCOMM’13 and SIGCOMM ICN’13 workshop*, extended abstract, Hong Kong (Aug 2013).
7. Jianqing Zhang, Qinghua Li, Eve M. Schooler, “iHEMS: An Information-Centric Approach to Secure Home Energy Management”, *IEEE 3rd International Conference on Smart Grid Communications*, SmartGridComm’12, Tainan City, Taiwan (Nov 2012).

Gartner Hype Curve

July 2015




Source: Gartner (July 2015)

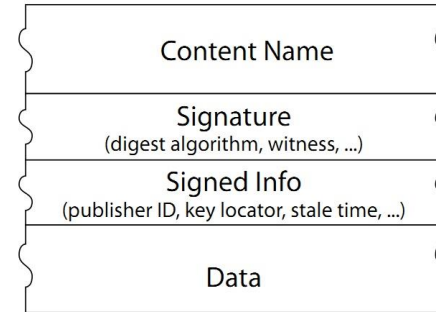
Attractive ICN Properties

- Name-based data routing
- Distributed data caching
- Self-contained data security

Interest packet



Data packet

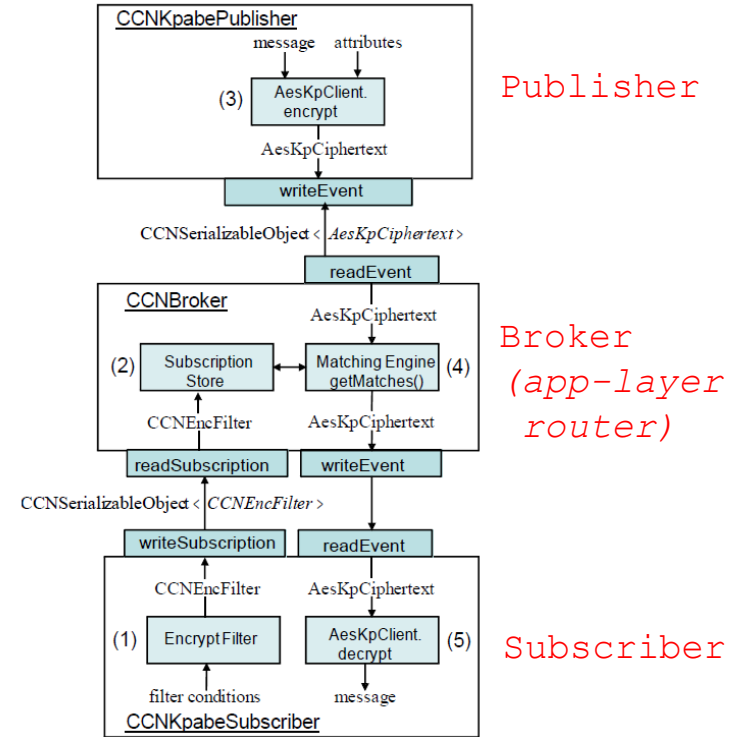


ICN Challenges


- Data & Interest confidentiality without secret key sharing
- Enforce fine-grained data access control policies in a distributed manner

Approach: Content-Centric Privacy

- Enrich ICN with pub-sub layer
 - data described by **attributes** vs strictly name
- Support **fine-grain** constraints on attributes
 - to describe data access control policies
 - e.g., *who/what has access to the data what conditions*
 - to express Interests
- Attach **access control** mechanisms
 - to the data & decryption keys




IoT Data Infographic: *Data Tidal Wave*



Growth of Cloud-Based Data as a Percentage of Total Data

- Total amount of data
- Data in the cloud




Source: IDC, The Digital Universe in 2020 (2012)

Why Critical to Solve?

2020 Expectations :

Video from huge #s of connected cameras:

- Forecast for a total of 129 yottabytes generated by 2020, of which 41% will come from sensors and 59% from cameras
(ABI Research, April 2015)
- 180/360-degree IP network cameras are the fastest growing product segment in video surveillance and IP panoramic network cameras are forecast to increase global unit shipments by more than 60% YoY
(IHS, 2015)

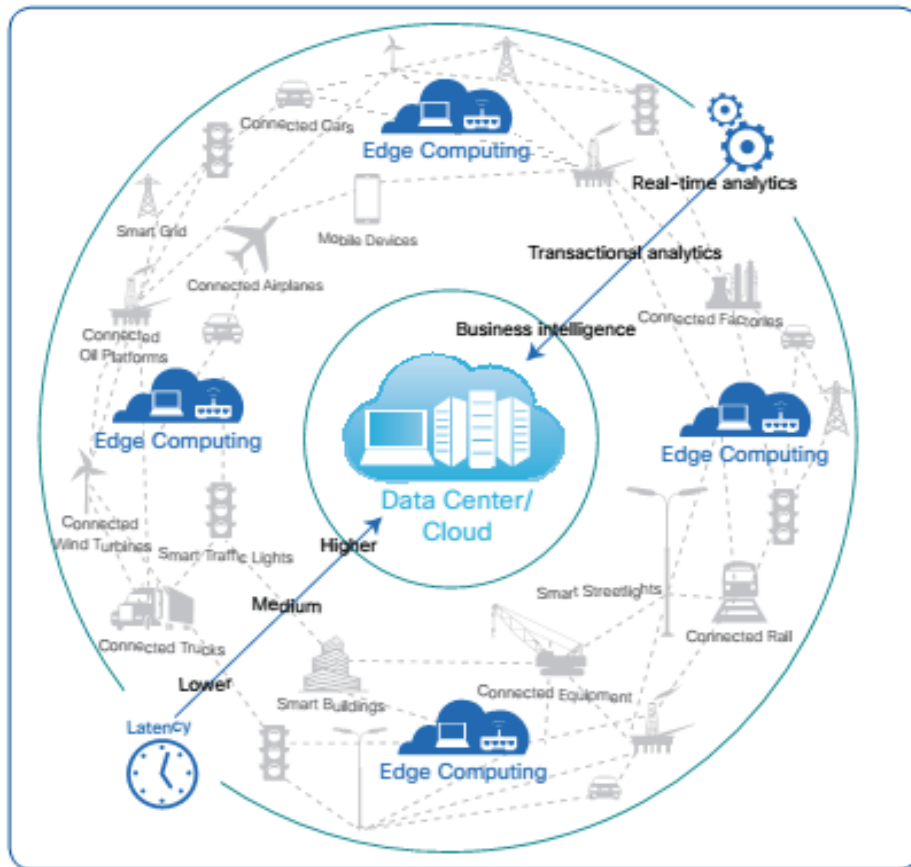


ABI Research

Why Critical to Solve?

2020 Expectations: Storage & Compute at the Edge

- By 2020, 40% of all data will come from IoT devices and sensors - nearly reaching 90% of the world's data created in last 2 years (*Cisco Consulting Services, 2014*)
- Edge computing ensures that the right processing takes place at the right time & the right place



Source: Cisco, 2014